Modelling Career Trajectories of Cricket Players Using Gaussian Processes

Oliver G. Stevenson and Brendon J. Brewer

Abstract In the sport of cricket, variations in a player's batting ability can usually be measured on one of two scales. Short-term changes in ability that are observed during a single innings, and long-term changes that are witnessed between matches, over entire playing careers. To measure long-term variations, we derive a Bayesian parametric model that uses a Gaussian process to measure and predict how the batting abilities of international cricketers fluctuate between innings. The model is fitted using nested sampling given its high dimensionality and for ease of model comparison. Generally speaking, the results support an anecdotal description of a typical sporting career. Young players tend to begin their careers with some raw ability, which improves over time as a result of coaching, experience and other external circumstances. Eventually, players reach the peak of their career, after which ability tends to decline. The model provides more accurate quantifications of current and future player batting abilities than traditional cricketing statistics, such as the batting average. The results allow us to identify which players are improving or deteriorating in terms of batting ability, which has practical implications in terms of player comparison, talent identification and team selection policy.

Key words: cricket, Gaussian processes, nested sampling

1 Introduction

As a sport, cricket is a statistician's dream. The game is steeped in numerous statistical and record-keeping traditions, with the first known recorded scorecards dating

───────────────

Oliver G. Stevenson
University of Auckland, e-mail: `o.stevenson@auckland.ac.nz`

Brendon J. Brewer
University of Auckland, e-mail: `bj.brewer@auckland.ac.nz`

as far back as 1776. Given the statistical culture that has developed with the growth of cricket, using numeric data to quantify individual players' abilities is not a new concept. However, despite the abundance of data available, cricket has only recently attracted the attention of statistical analysts in the public realm. This is potentially due to many previous academic studies being focused on the likes of achieving a fair result [8, 11, 6] and match outcome prediction [4, 1, 17, 3], rather than statistical applications that measure and predict individual player abilities and performances.

For as long as the modern game has existed, a player's batting ability has primarily been recognized using the *batting average*; in general, the higher the batting average, the better the player is at batting. However, the batting average fails to tell us about variations in a player's batting ability, which can usually be attributed to one of two scales: (1) short-term changes in ability that are observed *during* or *within* a single innings, due to factors such as adapting to the local pitch and weather conditions (commonly referred to as 'getting your eye in' within the cricketing community), and (2) long-term changes that are observed *between* innings, over entire playing careers, due to the likes of age, experience and fitness levels.

Early studies provided empirical evidence to support the claim that a batsman's score could be modelled using a geometric progression, suggesting players bat with a somewhat constant ability during an innings [10]. However, it has since been shown that the geometric assumptions do not hold for many players, due to the inflated number of scores of 0 that are present in many players' career records [12, 2]. Rather than model batting scores, [12] and [5] used nonparametric and parametric hazard functions respectively, to measure how dismissal probabilities change with a batsman's score. Estimating a batsman's hazard function, $H(x)$, which represents the probability of getting out on score x, allows us to observe how a player's ability varies over the course of an innings. Both studies found that batsmen appeared to be more likely to get out while on low scores – early in their innings – than on higher scores, supporting the idea of 'getting your eye in'.

In order to quantify the effects of 'getting your eye in', [16] proposed an alternative means of measuring how player ability varies during an innings. The authors use a Bayesian parametric model to estimate the hazard function, allowing for a smooth transition in estimated dismissal probabilities between scores, rather than the sudden, unrealistic jumps seen in [12] and to a lesser extent [5]. For the vast majority of past and present international Test players, [16] found overwhelming evidence to suggest that players perform with decreased batting abilities early in an innings and improve as they score runs, further supporting the notion of 'getting your eye in'.

1.1 Modelling Between-Innings Changes in Batting Ability

While there is plenty of evidence to suggest that players do not bat with some constant ability during an innings, it is also unlikely that a player bats with some constant ability throughout their entire career. Instead, variations in a player's underly-

ing ability are likely to occur *between* innings, due to factors such as how well the player has been performing recently (referred to as 'form' in cricket).

If batting form were to have a significant impact on player performance, we should be able to identify extended periods of players' careers with sequences of high scores (indicating the player was 'in' form) and sequences of low scores (indicating the player was 'out of' form). On the contrary, [9] found little empirical evidence to support this idea. Instead, for the majority of players analyzed in the study, the authors suggest that public perceptions of batting form tend to be overestimated, with many players' scores able to be modelled using a random sequence.

Within a Bayesian framework, [13] employed the use of a hidden Markov model to determine whether a batsman is in or out of form. The model estimates a number, K, of 'underlying batting states' for each player, including the expected number of runs to be scored when in each of the K states. Parameters that measure: *availability* (the probability a batsman is in form for a given match), *reliability* (the probability a batsman is in form for the next n matches) and *mean time to failure* (the expected number of innings a batsman will play before he is out of form), were also estimated for each batsman. However, a drawback of this approach is that the model requires an explicit specification of what constitutes an out of form state. The authors specify a batting state that has a posterior expected median number of runs scored of less than 25, as being out of form, and all other states as being in form. While in the context of one day or Twenty20 cricket this is not necessarily an unreasonable specification, there are numerous arguments that could be made to justify a low score, scored at a high strike rate, as a successful innings.

In this paper, we extend the Bayesian parametric model detailed in [16], such that we can not only measure and predict how player batting abilities fluctuate during an innings, but also between innings, over the course of entire playing careers. This allows us to treat batting form as continuous, rather than binary; instead of defining players as 'in' or 'out' of form, we can describe players as improving or deteriorating in terms of batting ability. At this stage our focus is on longer form Test and first-class cricket, as limited overs cricket introduces a number of match-specific complications [7].

2 Model Specification

The derivation of the model likelihood follows the method detailed in [16]. If $X \in \{0,1,2,3,...\}$ is the number of runs a batsman is currently on, we define a *hazard function*, $H(x) \in [0,1]$, as the probability a batsman gets out on score x. Assuming a functional form for $H(x)$, conditional on some parameters θ, we can calculate the probability distribution for X as follows:

$$P(X = x) = H(x) \prod_{a=0}^{x-1} [1 - H(a)]. \tag{1}$$

For any given value of x, this can be thought of as the probability of a batsman surviving up until score x, then being dismissed. However, in cricket there are a number of instances where a batsman's innings may end without being dismissed (referred to as a 'not out' score). Therefore, in the case of not out scores, we compute $P(X \geq x)$ as the likelihood, rather than $P(X = x)$. Comparable to right-censored observations in the context of survival analysis, this assumes that for not out scores the batsman would have gone on to score some unobserved score, conditional on their current score and their assumed hazard function.

Therefore, if I is the total number of innings a player has batted in and N is the number of not out scores, the probability distribution for a set of conditionally independent 'out' scores $\{x_i\}_{i=1}^{I-N}$ and 'not out' scores $\{y_i\}_{i=1}^{N}$ can be expressed as

$$p(\{x\}, \{y\}) = \prod_{i=1}^{I-N} \left(H(x_i) \prod_{a=0}^{x_i-1} [1 - H(a)] \right) \times \prod_{i=1}^{N} \left(\prod_{a=0}^{y_i-1} [1 - H(a)] \right). \tag{2}$$

When data $\{x, y\}$ are fixed and known, Equation 2 gives the likelihood for any proposed form of the hazard function, $H(x; \theta)$. Therefore, conditional on the set of parameters θ governing the form of $H(x)$, the log-likelihood is

$$\log \Big(L(\theta) \Big) = \sum_{i=1}^{I-N} \log H(x_i) + \sum_{i=1}^{I-N} \sum_{a=0}^{x_i-1} \log[1 - H(a)] + \sum_{i=1}^{N} \sum_{a=0}^{y_i-1} \log[1 - H(a)]. \tag{3}$$

2.1 Parameterizing the Hazard Function

The model likelihood in Equation 3 depends on the parameterization of the hazard function, $H(x)$. As per [16], we parameterize the hazard function in terms of an *effective average function*, $\mu(x)$, which represents a player's ability on score x, in terms of a batting average. Given the prevalence of the batting average in cricket, it is far more intuitive for players and coaches to think of ability in terms of batting averages, rather than dismissal probabilities. The hazard function can then be expressed in terms of the effective average function, $\mu(x)$, as follows

$$H(x) = \frac{1}{\mu(x) + 1} \tag{4}$$

where the effective average contains three parameters, $\theta = \{\mu_1, \mu_2, L\}$, and takes the following functional form

$$\mu(x) = \mu_2 + (\mu_1 - \mu_2) \exp\left(\frac{-x}{L} \right). \tag{5}$$

Here, μ_1 represents a player's initial batting ability when beginning a new innings, while μ_2 is the player's 'eye in' batting ability once used to the specific match conditions. Both μ_1 and μ_2 are expressed in terms of a batting average. The timescale parameter L, measures the speed of transition between μ_1 and μ_2 and is formally the e-folding time. By definition the e-folding time, L, signifies the number of runs scored for approximately 63% (formally $1 - \frac{1}{e}$) of the transition between μ_1 and μ_2 to take place and can be understood by analogy with a 'half-life'. This model specification allows us to answer questions about individual players, such as: (1) how well players perform when they first arrive at the crease, (2) how much better players perform once they have 'got their eye in' and (3) how long it takes them to 'get their eye in'.

2.2 Modelling Between-Innings Changes in Batting Ability

To extend the model further, such that we can measure variations in player batting ability between innings, we use the same likelihood function in Equation 3. However, we re-parameterize the effective average function to include a time component, t, such that

$$\mu(x,t) = \text{expected batting average on score } x, \text{ in } t^{th} \text{ career innings.} \quad (6)$$

For clarity, we will refer to $\mu(x)$ as the 'within-innings' effective average (explaining how ability changes within an innings). By marginalizing over all scores, x, we obtain the 'between-innings' effective average, $\nu(t)$, which explains how ability changes between innings, across a playing career.

$$\nu(t) = \text{expected batting average in } t^{th} \text{ career innings.} \quad (7)$$

When estimating $\nu(t)$, we need to account for variations in ability due to external factors such as: recent form, general improvements/deterioration in skill and the element of randomness associated with cricket. This is achieved by fitting a μ_2 parameter for each innings in a player's career, where μ_{2_t} represents a player's 'eye in' batting ability, corresponding to their t^{th} innings. We are then able to predict the expected batting average in each innings, $\nu(t)$, analytically using Equation 5.

To afford a player's underlying batting ability a reasonable amount of flexibility, the set of $\{\mu_{2_t}\}$ terms are modelled using a Gaussian process. A Gaussian process is fully specified by an underlying mean value, m, and covariance function, $K(X_i, X_j)$, which will determine by how much a player's batting ability can vary from innings to innings [14]. Our choice of covariance function is the commonly used squared exponential covariance, which contains scale and length parameters σ and ℓ.

Therefore, the model contains the set of parameters $\theta = \{\mu_1, \{\mu_{2_t}\}, L, m, \sigma, \ell\}$. The model structure with respect to parameters μ_1, L, C and D follows the model specification detailed in [16], with the parameters assigned the following prior distributions.

$$\mu_1 \leftarrow C\mu_2 \qquad\qquad \log(\mu_{2_t}) \sim \text{Gaussian process}(m, K(X_i, X_j; \sigma, \ell))$$
$$L \leftarrow D\mu_2 \qquad\qquad m \sim \text{Lognormal}(\log(25), 0.75^2)$$
$$C \sim \text{Beta}(1,2) \qquad\qquad \sigma \sim \text{Exponential}(10)$$
$$D \sim \text{Beta}(1,5) \qquad\qquad \ell \sim \text{Uniform}(0,100)$$

These priors are either non-informative or are relatively conservative, while loosely reflecting our cricketing knowledge. It is worth noting, that as we are measuring ability in terms of a batting average (which must be positive), we model $\log(\mu_{2_t})$, rather than just μ_{2_t}, to ensure positivity in our estimates.

As the model requires a set of $\{\mu_{2_t}\}$ parameters to be fitted (one for each innings played), the model can contain a large number of parameters for players who have enjoyed long international careers. Therefore, to fit the model we employ a C++ implementation of the nested sampling algorithm [15], which uses Metropolis-Hastings updates and is able to handle both high dimensional and multimodal problems that may arise. The model output provides us with the posterior distribution for each of the model parameters, as well as the marginal likelihood, which makes for trivial model comparison. For each player analyzed, we initialize the algorithm with 1000 particles and use 1000 MCMC steps per nested sampling iteration.

3 Analysis of Individual Players

3.1 Data

The data we use to fit the model are simply the Test career scores of an individual batsman and are obtained from Statsguru, the cricket statistics database on the Cricinfo website[1]. As the model assumes that a player's ability is not influenced by the specific match scenario, it is best suited to longer form cricket, such as Test matches, where there is generally minimal external pressure on batsmen to score runs at a prescribed rate.

3.2 Modelling Between-Innings Changes in Batting Ability

To illustrate the practical implications of the model, let us consider the Test match batting career of current New Zealand captain, Kane Williamson. The evolution of Williamson's between-innings effective average, $v(t)$, is shown in Figure 1 and suggests that early in his career, Williamson was not as good a batsman as he is today. In fact, it was not until playing in roughly 50 innings that he began to consistently bat *at least* as well as his current career average of 50.36. This is not surprising, as

---

[1] www.espncricinfo.com

it is a commonly held belief that many players need to play in a number of matches to 'find their feet' at the international level, before reaching their peak ability.
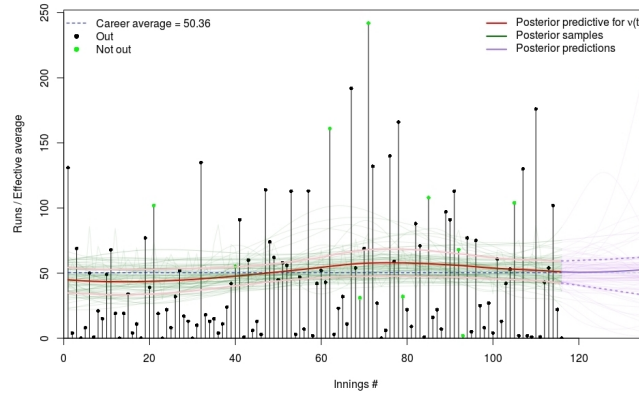

Fig. 1 Posterior predictive effective average function for $v(t)$ (red), fitted to Kane Williamson's Test career data, including a subset of posterior samples (green), future predictions (purple) and a 68% credible interval (pink/dotted purple).

To gain a better understanding of how Williamson compares to other batsmen globally, we can compare multiple players' effective average functions. Figure 2 compares the predictive effective average functions for the current top four batsmen in the world, as ranked by the official International Cricket Council (ICC) ratings[2].
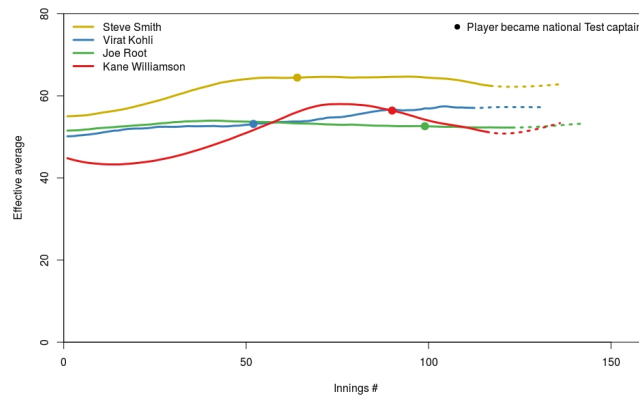

Fig. 2 Posterior predictive effective average functions, $v(t)$, for 'the big four', including predictions for the next 20 innings (dotted).

___

[2] As of 1[st] August, 2018: (1) Steve Smith, (2) Virat Kohli, (3) Joe Root and (4) Kane Williamson – commonly referred to as 'the big four'

As we might expect, all players appear to have improved in terms of batting ability since the start of their careers. Again, this supports the concept of 'finding your feet' at the international level, although different players appear to take different lengths of time to adjust to the demands of international cricket.

Table 1 shows each player's predicted effective average for their next innings, as well as their ICC rating. The order of these four players remains unchanged when ranking by predicted effective averages instead of ICC ratings, however, as we have computed the posterior predictive distributions for $v(t)$, our model has the added advantage of being able to quantify the differences in abilities between players. Rather than concluding 'Steve Smith is 26 rating points higher than Virat Kohli', we can make more useful statements such as: 'we expect Steve Smith to outscore Virat Kohli by 5.1 runs in their next respective innings' and 'Steve Smith has a 68.8% chance of outscoring Virat Kohli in their next respective innings'.

Table 1 Predicted effective averages, $v(t)$, for the next career innings for 'the big four'. The official ICC Test batting ratings (as of 1^{st} August, 2018) are shown for comparison.

Player	Career average	Predicted v(next innings)	ICC Rating (#)
S. Smith (AUS)	61.4	62.5	929 (1)
V. Kohli (IND)	53.4	57.4	903 (2)
J. Root (ENG)	52.6	52.6	855 (3)
K. Williamson (NZ)	50.4	51.2	847 (4)

4 Concluding Remarks and Future Work

We have presented a novel and more accurate method of quantifying player batting ability than traditional cricketing statistics, such as the batting average. The results provide support for the common cricketing belief of 'finding your feet', particularly for players beginning their international careers at a young age, with many batsmen taking a number of innings to reach their peak ability in the Test match arena. With respect to batting form, the model appears to reject the idea of recent performances as having a significant impact on innings in the near future. In particular, it appears that the effect of recent form varies greatly from player to player.

A major advantage of the model is that we are able to maintain an intuitive cricketing interpretation, allowing for the results and implications to be easily digested by coaches and selectors, who may have minimal statistical training. Additionally, we are able to make probabilistic statements and comparisons between players, allowing us to easily quantify differences in abilities and predict the real life impacts of selecting one player over another. As such, the findings have practical implications in terms of player comparison, talent identification, and team selection policy.

It is worth noting that we have ignored important variables, such as the number of balls faced in each innings, as well as the strength of the opposition. Currently, the model treats all runs scored equally. Implementing a means of incorporating more in-depth, ball-by-ball data and including the strength of opposition bowlers will reward players who consistently score highly against world-class bowling attacks.

References

[1] Bailey, M., Clarke, S.R.: Predicting the match outcome in one day international cricket matches, while the game is in progress. Journal of Sports Science & Medicine **5**(4), 480 (2006)

[2] Bracewell, P.J., Ruggiero, K.: A parametric control chart for monitoring individual batting performances in cricket. Journal of Quantitative Analysis in Sports **5**(3) (2009)

[3] Brooker, S., Hogan, S.: A method for inferring batting conditions in ODI cricket from historical data (2011)

[4] Brooks, R.D., Faff, R.W., Sokulsky, D.: An ordered response model of Test cricket performance. Applied Economics **34**(18), 2353–2365 (2002)

[5] Cai, T., Hyndman, R.J., Wand, M.: Mixed model-based hazard estimation. Journal of Computational and Graphical Statistics **11**(4), 784–798 (2002)

[6] Carter, M., Guthrie, G.: Cricket interruptus: fairness and incentive in limited overs cricket matches. Journal of the Operational Research Society **55**(8), 822–829 (2004)

[7] Davis, J., Perera, H., Swartz, T.B.: A simulator for Twenty20 cricket. Australian & New Zealand Journal of Statistics **57**(1), 55–71 (2015)

[8] Duckworth, F.C., Lewis, A.J.: A fair method for resetting the target in interrupted one-day cricket matches. Journal of the Operational Research Society **49**(3), 220–227 (1998)

[9] Durbach, I.N., Thiart, J.: On a common perception of a random sequence in cricket: application. South African Statistical Journal **41**(2), 161–187 (2007)

[10] Elderton, W., Wood, G.H.: Cricket scores and geometrical progression. Journal of the Royal Statistical Society **108**(1/2), 12–40 (1945)

[11] Ian, P., Thomas, J.: Rain rules for limited overs cricket and probabilities of victory. Journal of the Royal Statistical Society: Series D (The Statistician) **51**(2), 189–202 (2002)

[12] Kimber, A.C., Hansford, A.R.: A statistical analysis of batting in cricket. Journal of the Royal Statistical Society. Series A (Statistics in Society) pp. 443–455 (1993)

[13] Koulis, T., Muthukumarana, S., Briercliffe, C.D.: A Bayesian stochastic model for batting performance evaluation in one-day cricket. Journal of Quantitative Analysis in Sports **10**(1), 1–13 (2014)

[14] Rasmussen, C.E., Williams, C.K.I.: Gaussian Processes for Machine Learning. MIT Press (2006)

[15] Skilling, J.: Nested sampling for general Bayesian computation. Bayesian analysis **1**(4), 833–859 (2006)

[16] Stevenson, O.G., Brewer, B.J.: Bayesian survival analysis of batsmen in Test cricket. Journal of Quantitative Analysis in Sports **13**(1), 25–36 (2017)

[17] Swartz, T.B., Gill, P.S., Muthukumarana, S.: Modelling and simulation for one-day cricket. Canadian Journal of Statistics **37**(2), 143–160 (2009)